


NCCS Brown Bag Series



TotalView on Discover:


Part 2

ReplayEngine and MemoryScape

Chongxun (Doris) Pan

doris.pan@nasa.gov


June 14, 2012



Agenda For Part 1 (see previous tutorial)



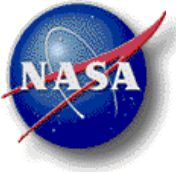
- Overview
- Basic Navigation and Control
- **Demo 1:** Basic Navigation and Control
 - ❖ Using an OpenMP Space weather application
- Debugging MPI Applications
- **Demo 2:** MPI Debugging:
 - ❖ Using GEOS5-AGCM



Agenda For Part 2




- Overview
- **Reverse Debugging with ReplayEngine**
- Demo 1: Replay Engine
 - ❖ Using an OpenMP Space weather application
- Memory Debugging with MemoryScape
- Demo 2: MemoryScape
 - ❖ Using GEOS AGCM



What is ReplayEngine ?




- Reverse Debugging
 - ❖ Let you step backward by function, line, or instruction
 - ❖ Capture and deterministically replay executions
- Major features
 - ❖ No recompilation
 - ❖ Designed to be used for parallel applications
- Why use ReplayEngine?
 - ❖ Eliminate restart cycle and hard-to-reproduce bugs



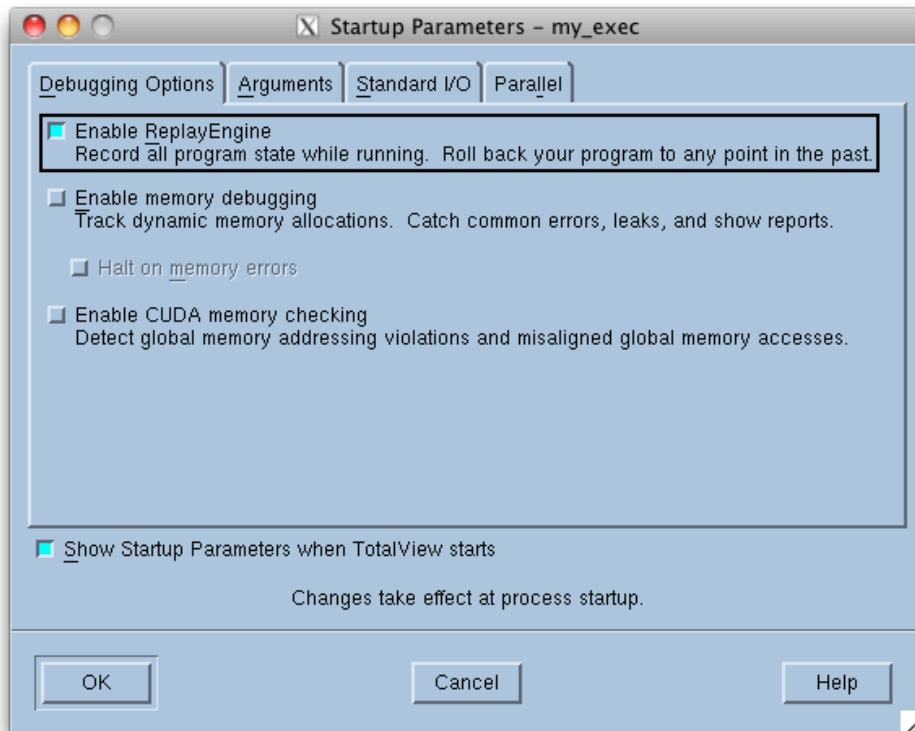
ReplayEngine



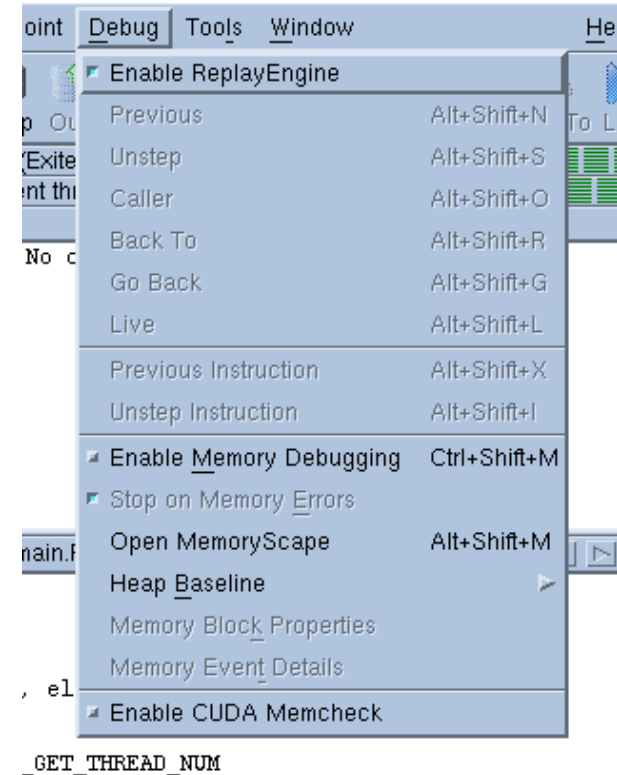
- **Record Mode** (saving execution history)
 - ❖ Capture the order of executions and changes to the data
- **Replay Mode** (when using a ReplayEngine command)
 - ❖ Like a “rewind” button on a DVR
- **Implications**
 - ❖ To ensure deterministic execution trajectory, it forces single thread execution at a time
 - ❖ During Replay mode, some operations (i.e., setting variables or full asynchronous thread control) are not available
 - ❖ Record mode could incur significant overhead (both in execution time and memory usage)




Enable ReplayEngine



1. Starting Totalview with
ReplayEngine enabled




2. After launching totalview, click the
“Enable ReplayEngine” Button.
ReplayEngine will be enabled after
“Restart”




ReplayEngine Usage Models



- **GoBack** (vs. **Go**): back to the last action point or the start of the recorded history
- **Prev** (vs. **Next**): back to previous statement executed. If the line has a function call, it skips over the call
- **Unstep** (vs. **Step**): back to previous statement executed. If the line has a function call, it moves to the last statement in the call
- **Caller** (vs. **Out**): back to before the current routine was called.







ReplayEngine Usage Models (Cont'd)

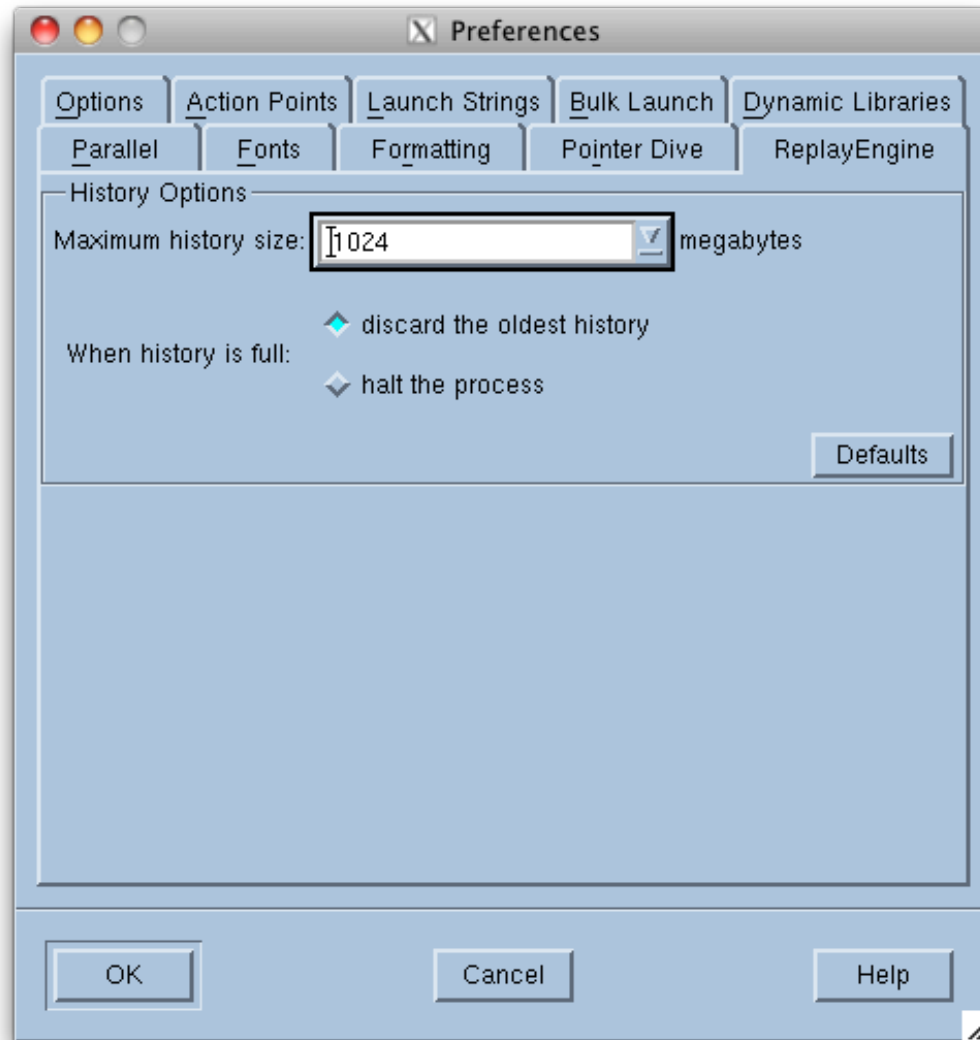


- **BackTo** (vs. **Run To**): back to the line you select
- **Live**: to shift from Replay mode to Record mode. It displays the statement that would have executed if you had not moved into Replay mode.






ReplayEngine Preferences



- Setting Maximum History size (default is unlimited)


- `setenv TVD_REPLAY_TMPDIR` to control the directory of saving the history information (default is /tmp)



Demo 1: ReplayEngine



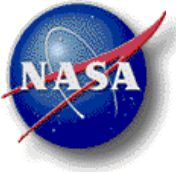
- Debugging an OpenMP Spaceweather application with ReplayEngine enabled
- Toolbar commands
- Preferences
- TVD_REPLAY_TMPDIR



Agenda For Part 2




- Overview
- Reverse Debugging with ReplayEngine
- Demo 1: Replay Engine
 - ❖ Using an OpenMP Space weather application
- **Memory Debugging with MemoryScape**
- Demo 2: MemoryScape
 - ❖ Using GEOS AGCM




Your Program data



- Your program data is organized in four segments:
 1. **Text (code):** machine code instructions
 2. **Data section:** static and global variables
 3. **Stack:** local (automatic) variables
 - ❖ Memory set aside for each thread of execution. It remains during program execution.
 - ❖ When a function is called, a block is reserved on the top of the stack for local (automatic) variables. It will go out of scope when the function returns and automatically deallocate.
 - ❖ Accessing stack in LIFO order – the most recently reserved block is the next block to be freed.
 - ❖ Easy to keep track of the stack -- Stack typically maps to the cache, so it is faster than the heap but smaller and expensive

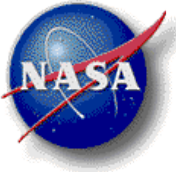




Your Program data -- What is **heap**?




4. **Heap**: all variables created or initialized *at* runtime are stored
 - ❖ **Dynamic memory allocations.** Heap size can grow as space is needed.
 - ❖ Variables created on the heap must be destroyed manually and never fall out of scope.
 - ❖ There's no enforced pattern to the allocation and deallocation of blocks from the heap; you can allocate a block at any time and free it at any time
 - ❖ Slower to allocate on the heap in comparison to variables on the stack
 - ❖ Can have fragmentation when there are a lot of allocations and deallocations
 - ❖ Responsible for memory leaks and many other memory bugs




What is MemoryScape?



- Runtime Memory Analysis to detect memory bugs
 - ❖ A memory bug is a mistake in heap memory usage
 - ❖ Failure to check for error conditions
 - ❖ Leaking: failure to free memory
 - ❖ Dangling references: failure to clear pointers
 - ❖ Memory corruption
 - ❖ Writing to memory not allocated
 - ❖ Over running array bounds
- Designed to be used with parallel applications.

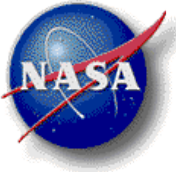




Why are memory bugs hard to find?




- Memory problems can be hidden
 - ❖ For a given scale or platform or problem, they may not be fatal
 - ❖ Failures could occur until modification, reuse of a component, or moving the application to a different cluster with a new OS
 - ❖ Libraries could be a source of problem
 - ❖ Can also manifest as “random crashes” or “random wrong answers”



MemoryScape Features:



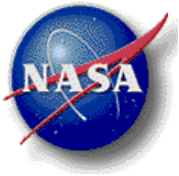
- Stopping execution when heap problems occur
- Viewing the heap graphically
- Leak detection and dangling pointer detection
- Memory Comparisons between processes
- Memory Corruption detection – Bound errors
 - ❖ **Guard Blocks** - allocated additional memory (8 bytes default) before or after a heap block. An event notification occurs when overwriting guard blocks.
 - ❖ **Red Zone** – allocated additional buffer before or after a heap block. An event notification occurs when either writing or reading to the Red Zone. (**High memory overhead!**)



Memory Debugger Features: (Cont'd)



- Block painting
 - ❖ Writing a bit pattern into allocated and deallocated blocks
 - ❖ Detecting the use of memory either before it is initialized or after it is deallocated
- Hoarding
 - ❖ Holding onto deallocated memory so it cannot be reused immediately
 - ❖ Not an often-used feature
 - ❖ Mostly used to detect problems related to memory being deallocated by one thread while another thread is using this memory



Start MemoryScape




1. Make sure your program is compiled with `-g`
2. `module load tool/tview-8.9.2-2`
3. `setenv PATH $PATH:/usr/local/toolworks/memryscape.3.2.2-2/bin/`
4. Launch totalview as usual, i.e.,
`totalview <executable>`

Then click “Enable memory debugging” in the “Startup Parameters” window

Or just type:


`memscape <executable>`



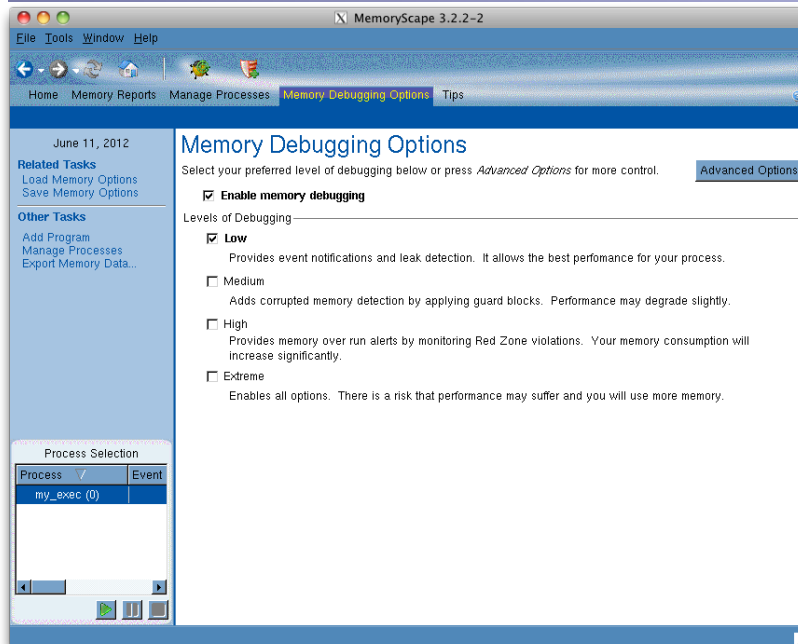
Strategies for Memory Debugging for Parallel Applications



- Run (or step through) the application and see if memory events are detected
- View memory usage across the MPI job
 - ❖ Compare memory footprints of the processes
- Gather heap information/Leak reports in all processes of the MPI job
 - ❖ Select and examine individually
 - ❖ Look at the allocation pattern
 - ❖ Look for leaks
 - ❖ Compare with the “diff” mechanism
 - ❖ Look for major differences



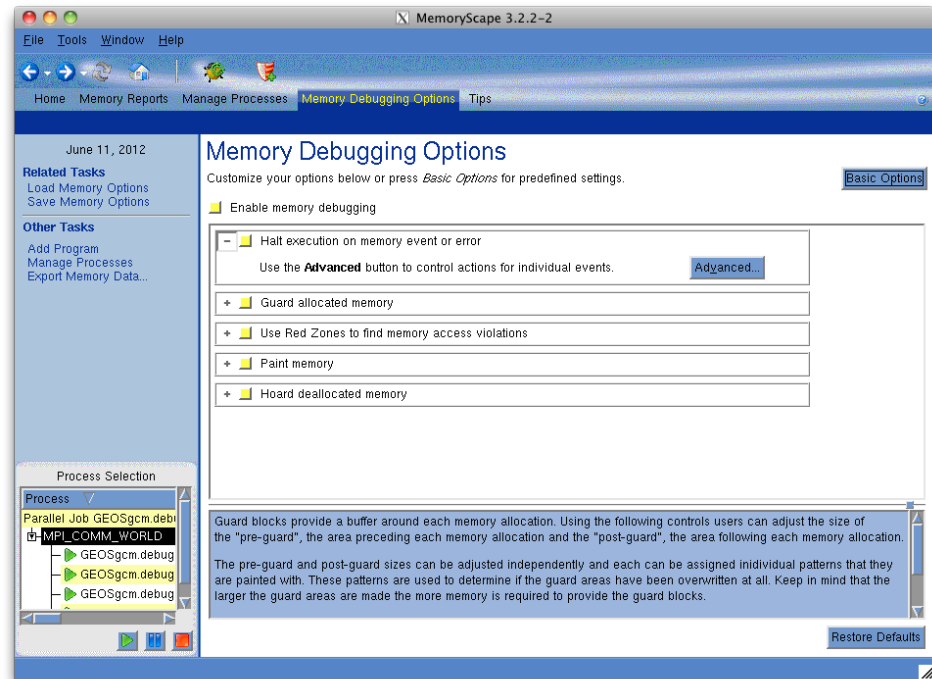
Memory Debugging Options




Select **Medium** only when you need to check for corrupted memory!

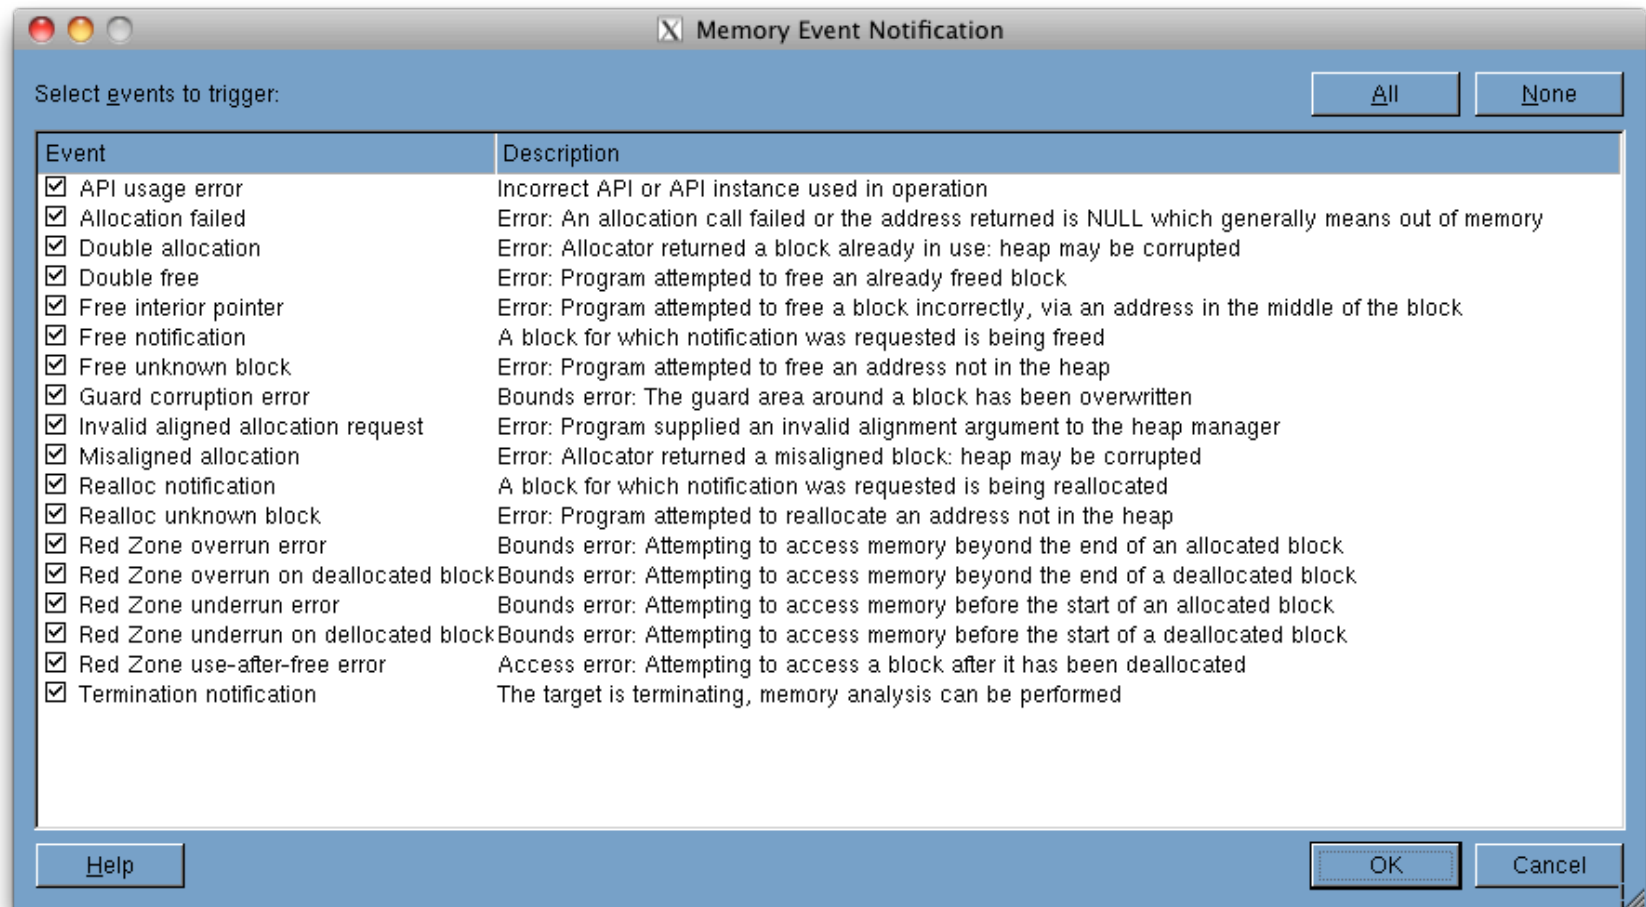
Select **High** only when you need to check for memory overruns!

Low is the default;
Medium adds guard blocks;
High adds Red Zone detection;
Extreme adds above all plus paint memory and hoarding.







What are memory events?



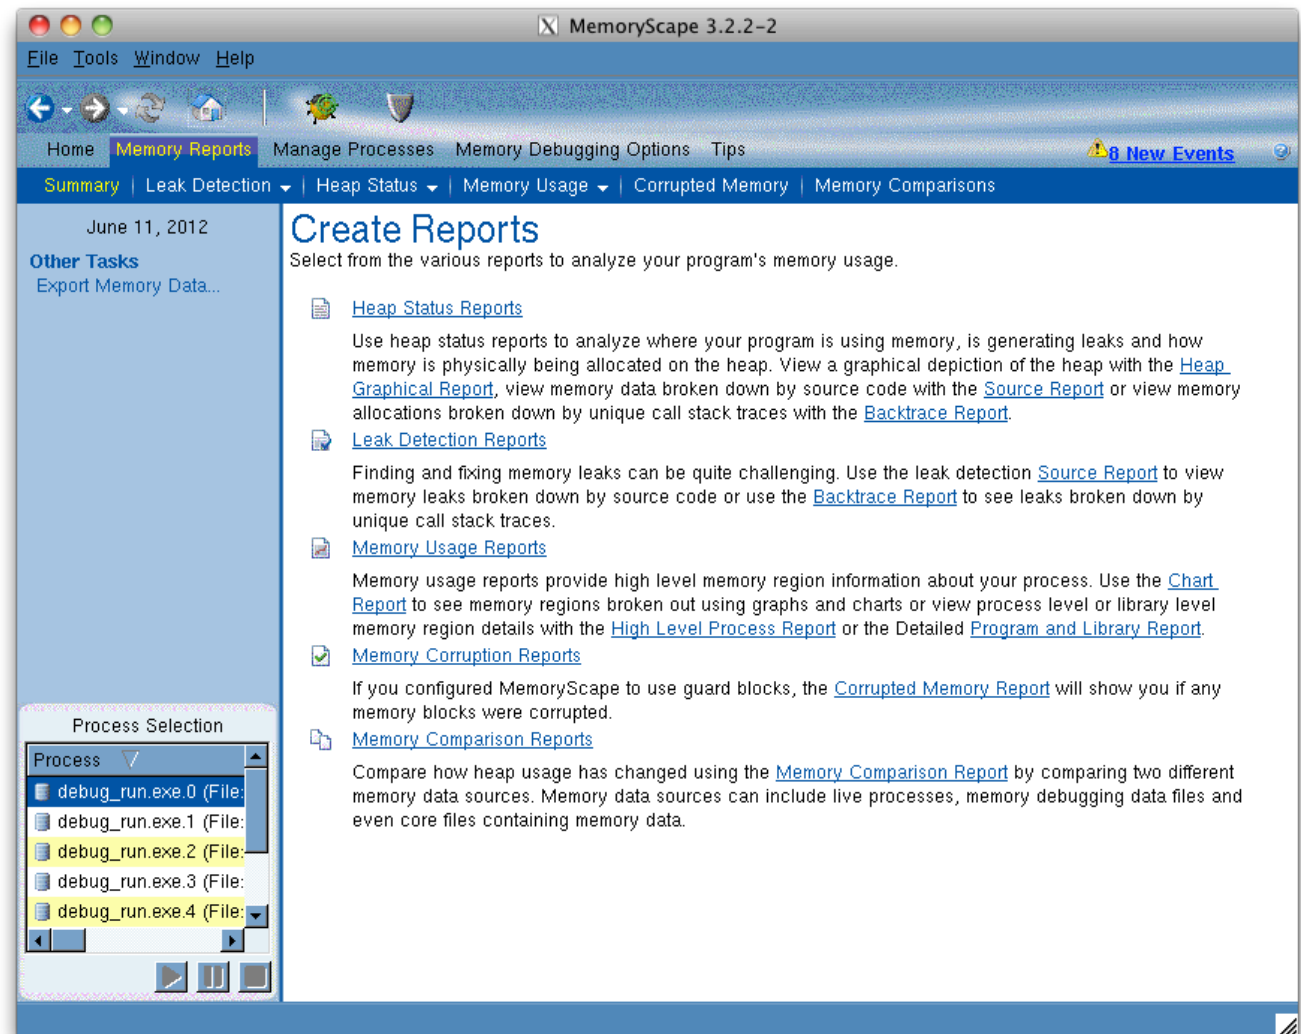
When one of these errors occurs, MemoryScape places event indicators by the process and at the top of the window

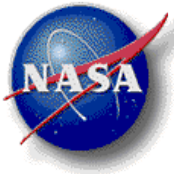


Analysis of memory debugging data

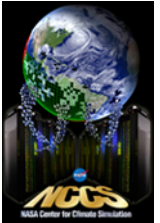


- **Heap status**
 - Graphical report
 - Source report
 - Backtrace report
- **Leak detection**
 - Source report
 - Backtrace report
- **Memory usage**
 - Chart report
 - High-level table
 - Detailed table
- **Memory corruption reports**
 - If guard blocks are chosen
- **Memory comparison among processes**





Heap Status Graphic Report

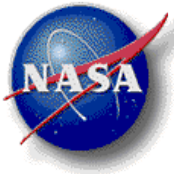


- Filtering reports is usually necessary to suppresses the display of too much information.

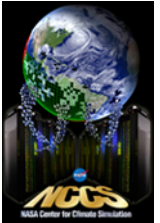


Process	Bytes	Count	Begin Address
Process 1 (10095): filterapp	1110.05KB	5531	
myClassB.cxx	1067.50KB	5404	
myClassA.cxx	28.00KB	56	
main.cxx	14.55KB	71	
main	14.55KB	71	
Line 17	14.00KB	28	
Line 19	336	28	
Line 20	168	14	

- Right-click routine name or line number in Heap Status Source or Backtrace report or in a Heap Status Graphical report, and choose “Filter out this entry”



Heap Status Source Report



MemoryScape 3.2.2-2

File Tools Window Help

Home **Memory Reports** Manage Processes Memory Debugging Options Tips

Summary | Leak Detection | **Heap Status** | Memory Usage | Corrupted Memory | Memory Comparisons

June 11, 2012

Save Data
Save Report...
Export Memory Data...

Heap Status Reports
Graphical Report
Backtrace Report

Other Reports Categories
Leak Detection Reports
Memory Usage Reports
Corrupted Memory Report
Compare Memory Usage

Other Tasks
Manage Filters

Process Selection

Process ▾
Parallel Job GEOSgcm.de
MPL_COMM_WORLD
GEOSgcm.debu
GEOSgcm.debu
GEOSgcm.debu

Heap Status Source Report

Data Source
◆ Allocations ◆ Deallocations ◆ Hoard ◆ Red Zones

Options
☐ Detect Leaks ☐ Relative to Baseline ☐ Enable Filtering

Process	Bytes	Count	Begin Address	End Address	Backtrace
Process 9: GEOSgcm.debug.x.1	4178.94MB	128747			
GEOSgcm.debug.x	4178.78MB	128448			
libdaploscm.so.2	128.00KB	1			
libdat2.so	33.55KB	296			
libibverbs.so.1	48	2			
Process 8: GEOSgcm.debug.x.0	4164.97MB	122960			
GEOSgcm.debug.x	4164.81MB	122661			
libdaploscm.so.2	128.00KB	1			
libdat2.so	33.55KB	296			
libibverbs.so.1	48	2			

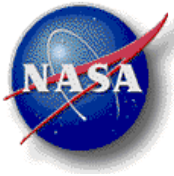
Backtrace

ID	Function	Line #	Source
5	malloc	166	malloc
	I_MPI_Collect_cpuinfo		libmpi:
	MPIR_Init_thread		libmpi:
	PMPI_Init_thread		libmpi:
	ZN5ESMC13VMK4initEi	15	m_iout
	_ZN5ESMC12VM10initializeEiPi	15	m_iout
	c_esmc_vminitialize_	15	m_iout
	esmf vmmod mp esmf vminitialize	15	m_iout

Source

na-2_5/src/GMAO_Shared/GMAO_mpeu/m_ioutil.F90

```
9 !  
10 ! m_ioutil is a module containing several port:  
11 ! some highly system dependent, but frequentl  
12 !  
13 ! INTERFACE:  
14  
15 module m_ioutil  
16 implicit none
```



Leak Detection Source Reports



MemoryScope 3.2.2-2

File Tools Window Help

Home Memory Reports Manage Processes Memory Debugging Options Tips

Summary | **Leak Detection** | Heap Status | Memory Usage | Corrupted Memory | Memory Comparisons

June 11, 2012

Save Data
Save Report...
Export Memory Data...

Leak Detection Reports
Backtrace Report

Other Reports Categories
Heap Status Reports
Memory Usage Reports
Corrupted Memory Report
Compare Memory Usage

Other
Manage Filters

Process Selection

Process: Parallel Job GEOSgcm.debug
MPI_COMM_WORLD
GEOSgcm.debug
GEOSgcm.debug
GEOSgcm.debug

Leak Detection Source Report

Options
☒ Relative to Baseline ☐ Enable Filtering

Process	Bytes	Count	Begin Address	End Address	Backtrace
Process 8: GEOSgcm.debug.x.0	9.84MB	1393			
GEOSgcm.debug.x	9.83MB	1232			
m_ioutil.F90	9.83MB	1232			
libdat2.so	11.58KB	161			
Process 7: GEOSgcm.debug.x.3	9.63MB	7329			
GEOSgcm.debug.x	9.62MB	7168			
m_ioutil.F90	9.62MB	7168			
libdat2.so	11.58KB	161			
Process 1: GEOSgcm.debug.x.2	9.42MB	7265			
GEOSgcm.debug.x	9.41MB	7104			

Backtrace

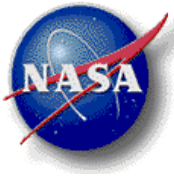
ID	Function	Line #	Source
9	PMPI_Init_thread	libmpi	
10	_ZN5ESMCi3VMK4initEi	15	m_ioutil
11	_ZN5ESMCi2VM10initializeEiPi	15	m_ioutil
12	c_esmc_vminitialize_	15	m_ioutil
13	esmf_vmmod_mp_esmf_vminitialize_	15	m_ioutil
14	esmf_initmod_mp_esmf_frameworkinternal...	15	m_ioutil
15	esmf_initmod_mp_esmf_initialize_	15	m_ioutil
16	MAPL_CAPMODmapl_cap	141	MAPL

Source


ja-2_5/src/GMAO_Shared/GMAO_mpeu/m_ioutil.F90

```
9 !  
10 ! m_ioutil is a module containing several por  
11 ! some highly system dependent, but frequen  
12 !  
13 ! INTERFACE:  
14  
15 module m_ioutil  
16 implicit none
```

Leak reports are essentially the same as the Heap status reports, only with fewer controls and less data.



Leak Detection Backtrace Reports



MemoryScape 3.2.2-2

File Tools Window Help

Home Memory Reports Manage Processes Memory Debugging Options Tips

Summary Leak Detection Heap Status Memory Usage Corrupted Memory Memory Comparisons

June 11, 2012

Save Data
Save Report...
Export Memory Data...

Leak Detection Reports
Source Report

Other Reports Categories
Heap Status Reports
Memory Usage Reports
Corrupted Memory Report
Compare Memory Usage

Other
Manage Filters

Process Selection

Process
Parallel Job GEOSgcm.debr
MPI_COMM_WORLD
GEOSgcm.debug
GEOSgcm.debug
GEOSgcm.debug

Leak Detection Backtrace Report

Options
☐ Relative to Baseline ☒ Enable Filtering


	Total Bytes	Count	Function	Line #	Source Information
gcm.debug.x.3	9.63MB	7329			
	13.12KB	1679			
			malloc	166	malloc_wrappers_dlopen.c
			for_allocate	15	m_ioutil.F90
			arrayscatter_r4_2	3016	MAPL_Comms__f90
			mapl_varread_r4_2d	2400	MAPL_IO__f90
			mapl_varread_r4_3d	2418	MAPL_IO__f90
			mapl_fieldread	1981	MAPL_IO__f90
			mapl_statevarread	1723	MAPL_IO__f90
			MAPL_GENERICMOD'mapl_esmfstateread...	4988	MAPL_Generic.F90
			MAPL_GENERICMOD'mapl_genericinitiali...	1241	MAPL_Generic.F90
			_ZN5ESMCI6FTable12callVFuncPtrEPKc...	15	m_ioutil.F90
			ESMCI_FTableCallEntryPointVMHop	15	m_ioutil.F90

Source

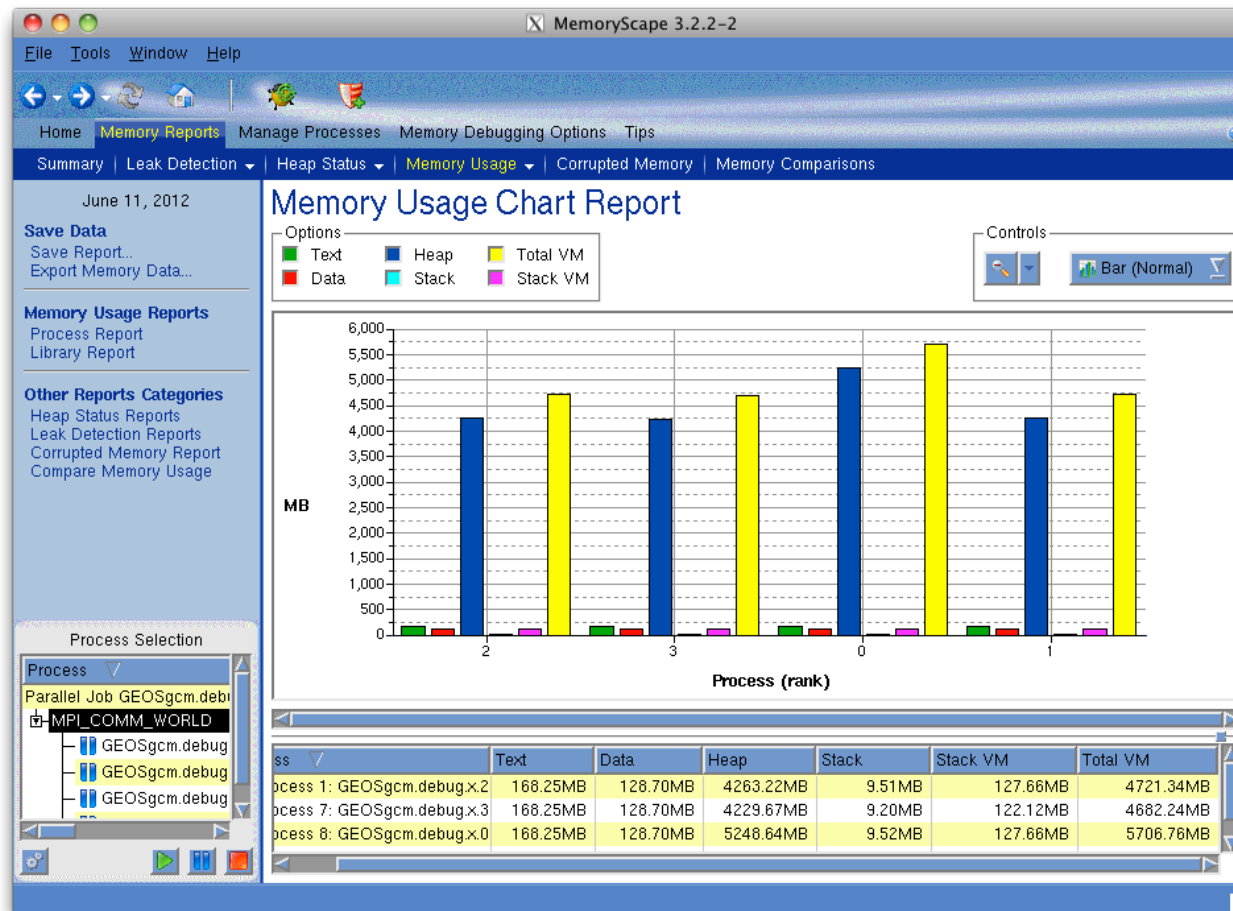
/gpfs/dnb31/cpan2/mybenchmark/GEOSgcm.Fortuna-2_5/src/GMAO_Shared/MAPL_Base/MAPL_Generic.F90

```
1238 LABEL="INTERNAL_HEADER:", &  
1239 RC=STATUS)  
1240 VERIFY_(STATUS)  
1241 call MAPL_ESMFStateReadFromFile(STATE%INTERNAL, CLOCK, FILENAME, &  
1242 FILETYPE, STATE, ndr=0, RC=STATUS)
```

Backtrace Report contains similar info as the source report. It is useful in helping you coordinate info in different screens and tabs as it does not change from report to report.




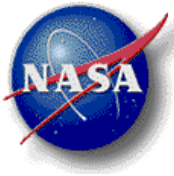
Memory Usage Chart Report




Differ slightly from the Heap Status reports:

- Memory usage data is obtained from the OS perspective, including overhead of the MemoryScape itself;
- Heap status is obtained from monitoring program requests and releases of memory





Memory Comparison Report



MemoryScape 3.2.2-2

File Tools Window Help

Home **Memory Reports** Manage Processes Memory Debugging Options Tips

Summary | Leak Detection | Heap Status | Memory Usage | Corrupted Memory | **Memory Comparisons**

June 11, 2012

Save View Options
Save Report...
Export Memory Data...

Other Reports Categories
Heap Status Reports
Memory Usage Reports
Corrupted Memory Report
Compare Memory Usage

Memory Comparison Report

Data Source: Allocations, Deallocations, Red Zones, Leaks, Hoard

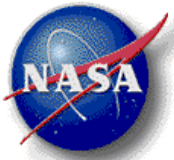
Process Comparisons:
Session 1: GEOSgcm.debug.x.1
Session 2: GEOSgcm.debug.x.0
Reverse Diff

Process	Bytes Session 2	Bytes Session 1	Bytes Difference	Count Session 2
GEOSgcm.debug.x.1/GEOSgcm.debug.x.0	413.49MB	427.46MB	-13.96MB	
GEOSgcm.debug.x	413.49MB	427.46MB	-13.96MB	
MAPL_CFIO.F90	676.00KB	338.00KB	338.00KB	
MAPL_Generic.F90	24.36KB	24.09KB	272	
MAPL_Comms__f90	0	338.00KB	-338.00KB	
m_ioutil.F90	412.81MB	426.77MB	-13.96MB	
for_get_vm	5.01MB	744	5.01 MB	
Line 15	5.01MB	744	5.01 MB	
_ZN5ESMCi8DistGrid14setArbSeqIndexE...	173.56KB	383.09KB	-209.53KB	
_ZN5ESMCi8DistGrid9constructEIIPIS1_S...	294.57KB	619.43KB	-324.86KB	
for_allocate	407.34MB	425.79MB	-18.45MB	


Session 1 Source: /gpfs/dnb31/cpan2/mybenchmark/GEOSgcm.Fortuna-2_5/src/GMAO_Shared/GMAO_mpeu/m_ioutil.F90

Session 2 Source: /gpfs/dnb31/cpan2/mybenchmark/GEOSgcm.Fortuna-2_5/src/GMAO_Shared/GMAO_mpeu/m_ioutil.F90

14
15 module m_ioutil



Memory Corruption Report



Memory Corruption Reports are generated only if you select “Medium” or higher debugging option

MemoryScape 2X.0.0-5

File Tools Window Help

Home Memory Reports Manage Processes Memory Debugging Options Tips

Summary Leak Detection Heap Status Memory Usage Corrupted Memory Memory Comparisons

June 11, 2009

Save Data
Save Report...
Export Memory Data...

Other Reports Categories
Heap Status Reports
Memory Usage Reports
Leak Detection Reports
Compare Memory Usage

Other Tasks
Manage Filters

Process Selection
Process /
RingBufferSPST (20352)

Corrupted Memory Report

Options
☐ Enable Filtering

	Preceding Block	Corrupted Block	Following Block
1	0x05000150 9 bytes 0x05000167	0x05000410 1624 bytes 0x050004ff	0x05000520
2	0x05000410 1624 bytes 0x050004ff	0x05000520 1624 bytes 0x0500052f	0x05000540
	0x05000520 1624 bytes 0x0500052f	0x05000540 1624 bytes 0x0500054f	0x05000560

Backtrace/Source | Memory Content

Process	Function	Line #	Source Information
Process	malloc	166	malloc_wrappers.dll
main		88	RingBufferSPST.c
__libc_start_main			libc.so.6
_start			RingBufferSPST

Source
e:\barry\tests\RingBuffer\RingBufferSPST.c


```
84 pCursor = pstructHead;  
85  
86 for(i=0; i<BLOCK_COUNT; i++)  
87 {  
88     pCursor->pBlock = (char*)m  
89     pCursor->pstructNext = /struct
```

- The top section graphically displaying each corruption.
- The bottom section contains a backtrace and the allocation source line.

Orange: corrupted guard block

Light green: uncorrupted guard block


Dark green: allocated data block



Agenda For Part 2



- Overview
- Reverse Debugging with ReplayEngine
- Demo 1: Replay Engine
 - ❖ Using an OpenMP Space weather application
- Memory Debugging with MemoryScape
- Demo 2: MemoryScape
 - ❖ Using GEOS AGCM




Demo 2



- `xsub -I -V -l select=1:ncpus=12:mpiprocs=12,walltime=1:00:00 -W group_list=<your_group>`

Once the compute node is available:

- `cd "$GEOS_EXPDIR"`
- `source $GEOSBIN/g5_modules`
- re-compile GCM with `BOPT=g`
- `mpdboot -n 1 -r ssh -f $PBS_NODEFILE`
- `module load tool/tview-8.9.2.2`
- `cp` and `link` necessary files ..(all in `gcm_run.j`)
- `totalview ./GEOSgcm.debug.x`



More info and references...



- MemoryScape_User_Guide.pdf
- ReplayEngine_Getting_Started_Guide.pdf

Under /usr/local/toolworks/totalview.8.9.2-2/doc/pdf

- “Memory Debugging Parallel Programs” tutorial offered by the Totalview in SC09

<http://www.crc.nd.edu/~rich/SC09/docs/tut120/tut120.pdf>